When we think about markets creating opportunities for interaction between buyers and sellers, there is an implicit network encoding the access between buyers and sellers. In fact, there are several ways to use networks to model buyer-seller interaction, and here we discuss some of them.

Bipartite Graphs and Perfect Matchings

First, consider the case in which not all buyers have access to all sellers. There could be several reasons for this lack of access – it could be informational (certain sellers and buyers are not aware of each other), or institutional (regulations or conventions prohibit certain sellers from transacting with certain buyers), or, as we will see later, each buyer could prioritize the sellers somehow, and only be interested in trading with the highest-priority seller(s).

One way to encode the pattern of potential interaction is as follows: we create a node for each seller and a node for each buyer, and we draw an edge between a seller and a buyer if they have the potential to engage in a transaction. We call such a graph a *bipartite graph*, since its nodes can be divided into two parts (in this case the sellers and buyers), in such a way that each edge has one endpoint in each part. Figure 1(a) gives some examples of bipartite graphs on sets of sellers and buyers; bipartite graphs are often drawn as in this figure, with the nodes in the two parts arranged in parallel columns.

One of the most basic questions about bipartite graphs concerns the notion of a *perfect matching*. Suppose we have an equal number of sellers and buyers. We can ask: is it possible to pair up the sellers and buyers in such a way that each pair has the opportunity to trade (i.e. the two nodes in each pair are joined by an edge)? We call such a pairing a *perfect matching* (since all nodes are perfectly matched together, with no one left out). For our discussion below, it will be more useful to think of a perfect matching in the following equivalent way: it's a set of edges with the property that each node is the endpoint of exactly one of them. (Such a set of edges specifies the desired pairing.) In Figure 1(b), we show a perfect matching in the bipartite graph from Figure 1(a), with the edges joining matched pairs indicated as darker lines.

So if a bipartite graph has a perfect matching, it's easy to demonstrate this: you just indicate the edges that form the perfect matching. But what if a bipartite graph has no perfect matching? What could you show someone to convince them that there isn't one?

At first glance, this is not clear; one naturally worries that the only way to convince someone that there is no perfect matching is to plow through all the possibilities and show that no pairing works. But in fact there is a clean way to demonstrate that no perfect matching exists, as follows. First, for any set of nodes A on one side of a bipartite graph, we say that a node on the other side is a *neighbor* of A if it has an edge to some node in A. We define the *neighbor set* of A, denoted $N(A)$, to be the collection of all neighbors of

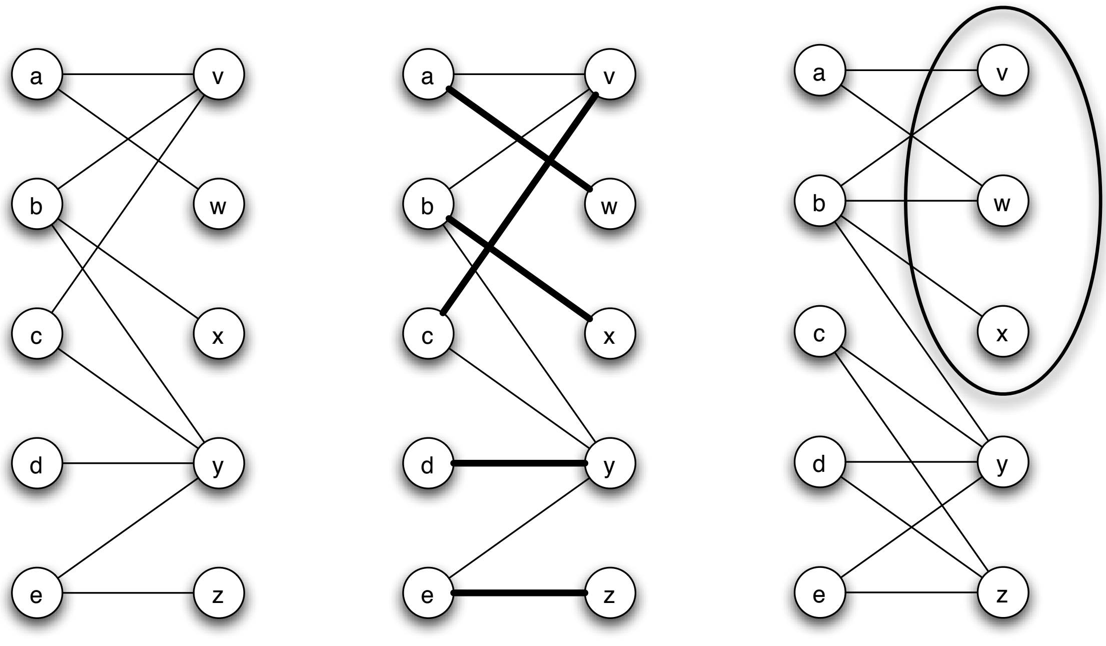

(a) *Bipartite Graph* (b) *A Perfect Matching* (c) *No Perfect Matching*

Figure 1: (a) An example of a bipartite graph, with sellers labeled a-e, buyers labeled v-z, and edges indicating opportunities for trade. (b) A perfect matching in this graph, indicated via the dark edges. (c) A different bipartite graph in which there is no perfect matching.

A. Finally, we say that a set A is *constricted* if A is strictly larger than $N(A)$ — that is, A contains strictly more nodes than $N(A)$ does.

Constricted sets in bipartite graphs are interesting objects; Figure 1(c) shows that the set A consisting of nodes v, w, and x is constricted, since its neighbor set consists of just a and b. If we find a constricted set A in a bipartite graph, then we can immediately conclude that the graph has no perfect matching: each node in A would have to be matched to a different node in $N(A)$, but there are more nodes in A than there are in $N(A)$, so this is not possible.

This much is not hard to see. What's also true, however, though far from obvious, is the following statement:

> *Claim: If a bipartite graph has no perfect matching, then it must contain a constricted set.*

This claim, known as the *Matching Theorem*, was independently discovered by Denes König in 1931 and Phillip Hall in 1935 [4]. Without the theorem, one might have imagined that a bipartite graph could fail to have a perfect matching for all sorts of reasons, some of them perhaps even too complicated to explain; but what the theorem says is that the simple notion of a constricted set is in fact the *only* obstacle to having a perfect matching. To put it another way, whenever a bipartite graph fails to have a perfect matching, there's an easy way to convince someone of this: just show them a constricted set.

In fact, given this Theorem, it is not hard to extend it to show something slightly stronger: if a bipartite graph has no perfect matching, then it has a constricted set on each of its two sides. (I.e., in our case, there is a constricted set of sellers *and* a constricted set of buyers; we've already seen that v, w, and x form a constricted set of buyers, and the nodes c, d, and e form a constricted set of sellers.) We will use this stronger fact later.

Valuations and Prices

The Matching Theorem is interesting from a graph-theoretic point of view, but the economic aspect is sublimated: it shows when we can pair off sellers and buyers, but it doesn't tell us anything about the payoffs in the resulting interaction. To get at these kinds of questions, we introduce some additional data into the problem.

Suppose as before that we have an equal number of sellers and buyers, and for now we will assume that all buyers have access to all sellers. However, each buyer j has her own *valuation* for each object being sold by seller i; we will denote this valuation by v_{ij}, since it depends on both i and j, and we will assume it is a non-negative whole number. One can imagine, for example, that each seller is trying to sell a house, and each buyer is looking to buy a house. Now, people's tastes in houses differ considerably: one might weight criteria based on size, location, age, or any of a number of other things. As a result, when buyer j sizes up the house being offered by seller i, she assigns a personal value of v_{ij} to it; this may differ from the value that some other buyer k places on the same house. Indeed, buyers j and k might not even agree on which of two houses is more valuable. (We assume that sellers, on the other hand, have a valuation of 0 for each house; they care only about receiving payment from buyers, which we will define shortly.)

Now, one question we might ask, given such data, is the welfare maximization question. In any perfect matching, each buyer receives a house that she values to some extent; which perfect matching maximizes the *total valuation*, the sum of each buyer's valuation for the house she receives? We're not going to address this question immediately, though notice for now that it at least subsumes the original perfect matching question that we discussed earlier: if all buyer valuations are either 1 (likes) or 0 (doesn't like), then finding a perfect matching with total valuation equal to n is the same as finding a perfect matching in the bipartite consisting just of the edges labeled 1.

The welfare maximization question assumes that you can simply announce who gets which house, and expect everyone to follow your instructions. But what if the choices made by buyers are determined in a more natural way, just by their own valuations and by prices offered by the sellers? What would happen then?

Prices. To explore these questions, let's suppose that each seller i offers to sell his house at a price $p_i \geq 0$. If a buyer j buys the house from seller i at this price, we will say that her *payoff* is her valuation for this house, minus the amount of money she had to pay: $v_{ij} - p_i$.

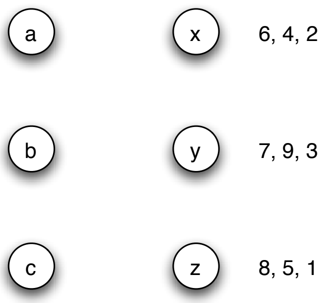

(a) *Buyer Valuations*

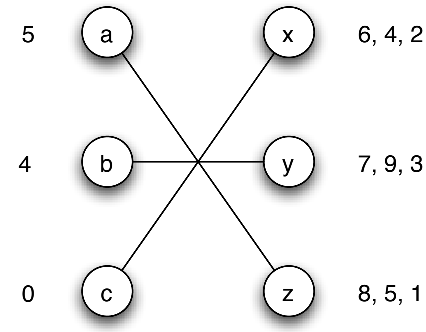

(b) *Market-Clearing Prices*

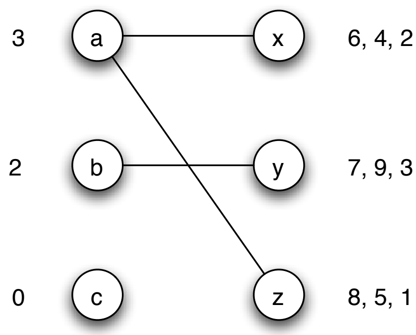

(c) *Prices that Don't Clear the Market*

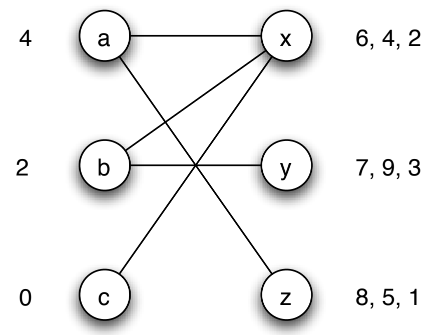

(d) *Market-Clearing Prices (Tie-Breaking Required)*

Figure 2: (a) Three sellers (a, b, and c) and three buyers (x, y, and z). For each buyer node, the valuations for the houses of the respective sellers appear in a list next to the node. (b) With prices 5, 4, and 0, each buyer creates a link to her preferred seller. The resulting set of edges is the preferred-seller graph for this set of prices. (c) The preferred seller graph for prices 3, 2, 0. (d) The preferred seller graph for prices 4, 2, 0.

So given a set of prices, if buyer j wants to maximize her payoff, she will buy from the seller i for which this quantity $v_{ij} - p_i$ is maximized — with the following caveats. First, if this quantity is maximized in a tie between several sellers, then the buyer can maximize her payoff by choosing any one of them. Second, if her payoff $v_{ij} - p_i$ is negative for every choice of seller i, then the buyer would prefer not to buy any house: we assume she can obtain a payoff of 0 by simply not transacting.

We will call the seller or sellers that maximize the payoff for buyer j the *preferred sellers* of buyer j, provided the payoff from these sellers is not negative. We say that buyer j has no preferred seller if the payoffs $v_{ij} - p_i$ are negative for all choices of i.

In Figures 2(b)-2(d), we show the results of three different sets of prices for the same set of buyer valuations. Note how the sets of preferred sellers for each buyer change depending on what the prices are. So for example, in Figure 2(b), buyer x would receive a payoff of

$6 - 5 = 1$ if she buys from a, a payoff of $4 - 4 = 0$ if she buys from b, and $2 - 0 = 2$ if she buys from c. This is why c is her unique preferred seller. One can similarly determine the payoffs for buyers y (2, 5, and 3) and z (3, 1, and 1) for transacting with sellers a, b, and c respectively.

Figure 2(b) has the particularly nice property that if each buyer simply seizes the house that she likes best, each buyer ends up with a different house: somehow the prices have perfectly resolved the contention for houses. (And this despite the fact that buyers x and z each valued the house of seller a the highest; it was the high price of 5 that dissuaded buyer x from pursuing it.)

We will call such a set of prices *market-clearing*, since they cause each house to get bought by a different buyer. Figure 2(c) shows an example of prices that are not market-clearing, since buyers x and z both want the house offered by seller a. And Figure 2(d) shows an example where one of the buyers (x) has several houses that tie for maximizing her payoff. If the buyers are able to coordinate so that x takes the house offered by the "right" one of her preferred sellers (c), then everyone gets a different house. As a result, we will say that this set of prices is market-clearing as well, even though a bit of coordination is required, based on the ties in the maximum payoffs, to get the market to clear. Ties like this may be inevitable: for example, if all buyers have the same valuations for everything, then no choice of prices will break this symmetry.

Given the possibility of ties, we will think about market-clearing prices more generally as follows. For a set of prices, we define the *preferred-seller graph* on buyers and sellers by simply constructing an edge between each buyer and her preferred seller or sellers. (There will be no edge out of a buyer if she has no preferred seller.) So in fact, Figures 2(b)-2(d) are just drawings of preferred-seller graphs for each of the three indicated sets of prices. Now we simply say: a set of prices is *market-clearing* if the resulting preferred-seller graph has a perfect matching.

Properties of Market-Clearing Prices. In a way, market-clearing prices feel a bit too good to be true: if sellers set prices the right way, then self-interest runs its course and (potentially with a bit of coordination over tie-breaking), all the buyers get out of each other's way and claim different houses. We've seen that such prices can be achieved in one very small example; but in fact, something much more general is true:

> *Existence of Market-Clearing Prices: For any set of buyer valuations, there exists a set of market-clearing prices.*

So market-clearing prices are not just a fortuitous outcome in certain cases; they are always present. This is far from obvious, and we will turn shortly to a method for constructing market-clearing prices that, in the process, proves they always exist.

Before doing this, we mention another natural question: the relationship between market-clearing prices and social welfare. Just because market-clearing prices cause all buyers to resolve their contention and get different houses, does this mean that the total valuation of the resulting matching will be good? In fact, there is something very strong that can be said here as well: market-clearing prices (for this buyer-seller matching problem) always provide socially optimal outcomes:

> *Optimality of Market-Clearing Prices: For any set of market-clearing prices, a perfect matching in the resulting preferred-seller graph has the maximum total valuation of any perfect matching of buyers and sellers.*

Compared with the previous claim on existence of market-clearing prices, this fact about optimality can be justified by a much shorter, if somewhat subtle, argument. This argument isn't crucial for anything that follows, so it can be skipped; but we describe it here because it is instructive.

The argument is as follows. Consider a set of market-clearing prices, and let M be a perfect matching in the preferred-seller graph. Since each buyer is grabbing a house that maximizes her payoff individually, M has the maximum total *payoff* to buyers of any perfect matching. Now how does total payoff relate to total valuation, which is what we're hoping that M maximizes? If buyer j chooses house i, then her valuation is v_{ij} and her payoff is $v_{ij} - p_i$. Thus, the total payoff to all buyers is simply the total valuation, minus the sum of all prices:

$$\text{Total Payoff of } M = \text{Total Valuation of } M - \text{Sum of all prices.}$$

But the sum of all prices is something that doesn't depend on which matching we choose (it's just the sum of everything the sellers are asking for, regardless of how they get paired up with buyers). So a matching M that maximizes the total payoff is also one that maximizes the total valuation.

Constructing a Set of Market-Clearing Prices

Now let's turn to the deeper challenge: understanding why market-clearing prices must always exist. We're going to do this by taking an arbitrary set of buyer valuations, and describing a procedure that arrives at market-clearing prices. The procedure will in fact be a kind of auction — not a single-item auction of the type we discussed earlier, but a more general kind taking into account the fact that there are multiple things being auctioned, and multiple buyers with different valuations. This particular auction procedure was described by the economists Demange, Gale, and Sotomayor in 1986 [2], but it's actually equivalent to a construction of market-clearing prices discovered by the Hungarian mathematician Egerváry seventy years earlier, in 1916 [4].

Here's how the auction works. Initially all sellers set their prices to 0. Buyers react by choosing their preferred seller(s), and we look at the resulting preferred-seller graph. If this graph has a perfect matching we're done. Otherwise — and this is the key point — there is a constricted set of buyers A. Consider the set of neighbors $N(A)$, which is a set of sellers. The buyers in A only want what the sellers in $N(A)$ have to sell, but there are fewer sellers in $N(A)$ than there are buyers in A. So the sellers in $N(A)$ are in "high demand" — too many buyers are interested in them. They respond by each raising their prices by one unit, and the auction then continues.

There's one more ingredient, which is a *balancing* operation on the prices. It will be useful to have our prices scaled so that the smallest one is 0. Thus, if we reach ever reach a point where all prices are strictly greater than 0 — suppose the smallest price has value $p > 0$ — then we "balance" the prices by substracting p from each one. This drops the lowest price to 0, and shifts all other prices by the same relative amount.

A general round of the auction looks like what we've just described.

(i) At the start of each round, there is a current set of prices, with the smallest one equal to 0.

(ii) We construct the preferred-seller graph and check whethere there is a perfect matching.

(iii) If there is, we're done: the current prices are market-clearing.

(iv) If not, we find a constricted set of buyers A and their neighbors $N(A)$.

(v) Each seller in $N(A)$ raises his price by one unit.

(vi) If necessary, we balance the prices — the same amount is subtracted from each price so that the smallest price becomes zero.

(vii) We now begin the next round of the auction, using these new prices.

Now, if we can simply show that the auction must eventually end — i.e. that the rounds cannot go on forever — then we're done, since the only way it can come to an end is with a set of market-clearing prices. The way we're going to show this is by identifying a precise sense in which a certain kind of "energy" is draining out of the auction as it runs; since the auction starts with only a bounded supply of this "energy" at the beginning, it must eventually run out.

Why the Auction Must Come to an End. Here is how we define the energy for use in this argument. For any current set of prices, define the *potential* of a buyer to be the maximum payoff she can currently get from any seller. (If the buyer has a preferred seller, this maximum is just the payoff from the preferred seller — but the potential can also be defined for buyers with no preferred sellers; it is simply a negative number in this case.) We also define the *potential* of a seller to be the current price he is charging. Finally, we define the *energy* of the auction to be the total potential of all participants, both buyers and sellers.

How does the energy of the auction behave as we run it? It begins at with all sellers having potential 0, and each buyer having a potential equal to her maximum valuation for any house — so the total potential at the start is some whole number $E_0 \geq 0$. Another thing to notice is that at the start of each round of the auction, everyone has potential at least 0. The sellers always have potential at least 0 since the prices are always at least 0. Because of balancing, the lowest price is always 0, and therefore each buyer is always doing at least as well as the option of buying a 0-cost item, which gives a payoff of at least 0. (This also

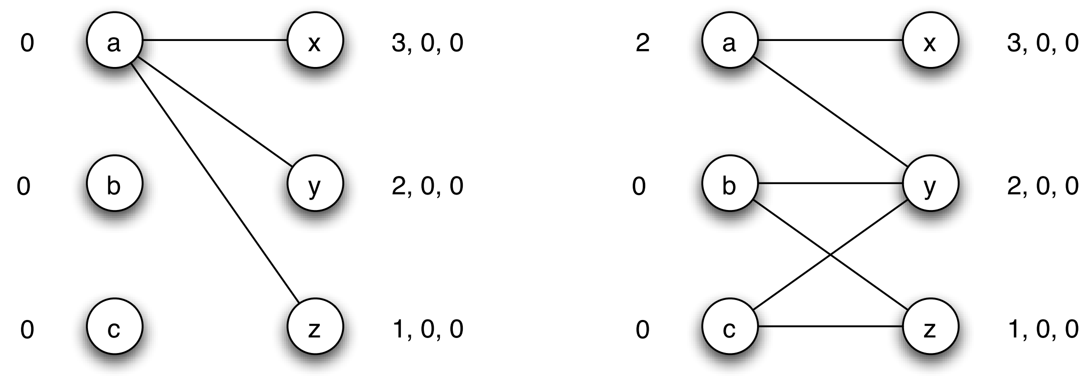

(a) *Start of the Auction* (b) *End of the Auction*

Figure 3: A single-item auction can be represented by the bipartite graph model: the item is represented by one seller node, and then there are additional seller nodes for which all buyers have 0 valuation. (a) The start of the bipartite graph auction. (b) The end of the bipartite graph auction, when buyer x gets the item at the valuation of buyer y.

means that each buyer has at least one preferred seller at the start of each round.) Finally, since the potentials are all at least 0 at the start of each round, so is the energy.

Now, the energy only changes when the prices change, and this only happens in steps (v) and (vi). Notice that the balancing of prices, as defined above, does not change the energy: if we subtract p from each price, then the potential of each seller drops by p, but the potential of each buyer goes up by p — it all cancels out. Finally, what happens to the energy in step (v) when the sellers in $N(A)$ all raise their prices by one unit? Each of their potentials goes up by one unit. But the potential of each buyer in A goes down by one unit, since all their preferred houses just got more expensive. Since A has strictly more nodes than $N(A)$ does, this means that the energy goes down by at least one unit more than it goes up, so it strictly decreases by at least one unit.

So what we've shown is that in each step that the auction runs, the energy decreases by at least one unit. It starts at some fixed value E_0, and it can't drop below 0, so the auction must come to an end within E_0 steps — and when it comes to an end, we have our market-clearing prices.

How Does this Relate to Single-Item Auctions?

Since we talked earlier about single-item auctions, and we now have this more complex type of auction based on bipartite graphs, it makes sense to ask how they relate to each other. In fact, there is a very natural way to view the single-item auction — both the outcome and the procedure itself — as a special case of the bipartite graph auction we've just defined. We can do this as follows.

Suppose we have a set of n buyers and a single seller auctioning an item; let buyer i

have valuation v_i for the item. To map this to our model based on perfect matchings, we need an equal number of buyers and sellers, but this is easily dealt with: we create $n-1$ "fake" additional sellers (who conceptually represent $n-1$ different ways to fail to acquire the item); we give buyer i a valuation of 0 for the item offered by each of these fake sellers. With the real seller labeled 1, this means we have $v_{i1} = v_i$, the valuation of buyer i for the real item; and $v_{ij} = 0$ for larger values of j.

Now we have a genuine instance of our bipartite graph model: from a perfect matching of buyers to sellers, we can see which buyer ends up paired with the real seller (this is the buyer who gets the item), and from a set of market-clearing prices, we will see what the real item sells for.

Moreover, the price-raising procedure to produce market-clearing prices — based on finding constricted sets — has a natural meaning here as well. The execution of the procedure on a simple example is shown in Figure 3. Initially, observe that all buyers will identify the real seller as their preferred seller (assuming that they all have positive valuations for the item). The first constricted set A we find is the set of all buyers, and $N(A)$ is just the single real seller. Thus, the seller raises his price by one unit. This continues as long as at least two buyers have the real seller as their unique preferred seller: they form a constricted set A with $N(A)$ equal to the real seller, and this seller raises his price by a unit. (Note that the prices of the fake items remain fixed at 0, so we never have to balance the prices.) Finally, when all but one buyer has identified other sellers as preferred sellers, the graph has a perfect matching. This happens at precisely the moment that the buyer with the second-highest valuation drops out — in other words, the buyer with the highest valuation gets the item, and pays the second-highest valuation. So the bipartite graph procedure precisely implements an ascending bid (English) auction.

Truthfulness

A major issue in our discussion of single-item auctions was that of truth-telling: we found that bidding one's true value is a dominant strategy in a second-price sealed-bid auction. In our bipartite graph model, we have thus far focused just on finding market-clearing prices given valuations, but we have not yet worried about whether it is a dominant strategy for buyers to actually reveal their true valuations.

In fact, a version of our bipartite auction will promote truth-telling; this was discovered independently by Demange and Leonard in 1983 [1, 3]. We describe their findings here briefly, without going into the details behind them; the specifics can be found in their papers, as well as in a subsequent book by Roth and Sotomayor that covers this area [5]. The result we describe here takes a little work to state precisely, but it has a very nice interpretation.

Among all possible sets of market-clearing prices, Demange and Leonard considered the one for which the sum of all the prices is as small as possible. Let M be the pairs in the (socially optimal) perfect matching associated with this set of prices. Let p_i be the price offered by seller i in this set, and suppose that in M, buyer j is paired with i. Demange and Leonard gave the following formula for the value of p_j, as follows. First, let a_{ij} be the total valuation of the matching M after the pair consisting of i and j is removed from it. Let b_{ij}

be the maximum valuation possible in a set of house purchases if we delete buyer j but not seller i. Notice that b_{ij} is at least as large as a_{ij}, since the remaining buyers could always ignore seller i if they wanted, thereby obtaining a_{ij}. What Demange and Leonard showed is that p_i is precisely the difference between these two quantities: $p_i = b_{ij} - a_{ij}$.

This can be viewed in the following suggestive way: when buyer j gets seller i's house, the remaining buyers can achieve a total valuation of a_{ij}. But without buyer j present, seller i becomes an option for everyone else, and so the remaining buyers can potentially do better, achieving b_{ij}. This difference $b_{ij} - a_{ij}$ in the total valuation obtainable by everyone else is exactly the price p_i that buyer j pays.

For our purposes here, the crucial point is that the price buyer j pays, namely $b_{ij} - a_{ij}$, does not depend at all on the valuations that buyer j has for the houses — it only depends on the valuations expressed by other buyers. This is the key to showing that buyer j has no incentive to misrepresent her true valuations; in other words, it can be shown that with these market-clearing prices, revealing one's true valuations is a dominant strategy. Again, one can look at the papers of Demange and Leonard [1, 3], or the book by Roth and Sotomayor [5], for more details about this.

References

[1] Gabrielle Demange. Strategyproofness in the assignment market game, 1982. Laboratiore d'Econometrie de l'Ecole Polytechnique.

[2] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[3] Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

[4] László Lovász and Michael Plummer. *Matching Theory*. North-Holland, 1986.

[5] Alvin E. Roth and Marilda Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.